On Formalisms for Dynamic Reconfiguration of Dependable Systems

Anirban Bhattacharyya, Andrey Mokhov, Ken Pierce
Newcastle University, UK

Manuel Mazzara
Politecnico di Milano, Italy

arXiv:1405.0319v2 [cs.SE] 7 May 2014

***Abstract*—Three formalisms of different kinds – VDM, Maude, and basic $\mathrm{CCS}^{\mathrm{dp}}$ – are evaluated for their suitability for the modelling and verification of dynamic software reconfiguration using as a case study the dynamic reconfiguration of a simple office workflow for order processing. The research is ongoing, and initial results are reported.**

I. State of the Art

We define the dynamic software reconfiguration of a system to be the change at runtime of the system's software components or their communication links, or the hardware location of the software components or their communication links. Dynamic software reconfiguration can help to increase the reliability of a system, and thereby increase its dependability. However, predicting the behaviour of the system during its dynamic reconfiguration is still a challenge [3], to which our research is a response. Existing research in dynamic software reconfiguration can be grouped into three cases:

1) The reconfiguration is near-instantaneous and its duration is negligible in comparison to the durations of application tasks. Executing tasks in the old configuration that are not in the new configuration are aborted. This is the traditional method of software reconfiguration, and it is applicable to small, simple systems running on a uniprocessor.
2) The duration of the reconfiguration interval is significant in comparison to the durations of application tasks, and any executing task in the old configuration that can interfere with a reconfiguration task is either aborted or suspended until the reconfiguration is complete. This is the most common method of software reconfiguration (see [17], [2], [4], and [10]), and it is applicable to some large, complex, distributed systems.
3) The duration of the reconfiguration interval is significant in comparison to the durations of application tasks, and tasks in the old configuration execute concurrently with reconfiguration tasks. This method avoids aborting tasks and reduces the delay on the application due to reconfiguration, but it introduces the possibility of functional and temporal interference between application and reconfiguration tasks. If the interference can be controlled, then this method is the most suitable for large, complex, distributed systems, including hard real-time systems, but it is also the least researched method of software reconfiguration. Existing research has focused on temporal interference between application and reconfiguration tasks, and on achieving schedulability guarantees (see [16], [18], [8], [15], [14], [7], and [9]). There is little research on formal verification of functional correctness in the presence of functional interference between application and reconfiguration tasks (see [13] and [5]).

II. Case study

We use as a case study the dynamic reconfiguration of a simple office workflow for order processing, commonly found in large and medium-sized organizations [6]. These workflows typically handle large numbers of orders. Furthermore, the organizational environment of a workflow can change in structure, procedures, policies, and legal obligations in a manner unforeseen by the original designers of the workflow. Therefore, it is necessary to support the unplanned change of these workflows. Furthermore, the state of an order in the old configuration may not correspond to any state of the order in the new configuration. These factors, taken in combination, imply that instantaneous reconfiguration of a workflow is not always possible; neither is it practical to delay or abort large numbers of orders because the workflow is being reconfigured. The only other possibility is to allow overlapping modes for the workflow during its reconfiguration.

The workflow consists of a graph of activities and the configuration is the structure of this graph. Initially, the workflow executes in Configuration 1 and must satisfy the requirements on Configuration 1. Subsequently, the workflow must be reconfigured through a process to Configuration 2 in such a way that the requirements on the reconfiguration process and on Configuration 2 are satisfied. The full details of the two configurations and their respective requirements are given in [11].

In order to achieve a smooth transition from Configuration 1 to Configuration 2 of the workflow, the process of reconfiguration must meet the following requirements:

1) Reconfiguration of a workflow should not necessarily result in the rejection of an order.
2) Any order being processed that was accepted **before** the start of the reconfiguration must satisfy all the requirements on Configuration 1.
3) Any order accepted **after** the start of the reconfiguration must satisfy all the requirements on Configuration 2.
4) The reconfiguration process must terminate.

III. Comparison of Formalisms

As discussed in previous work (see [12] and [1]), research shows that no single existing formalism is ideal for representing dynamic reconfiguration [19]. Therefore, we are using a portfolio of three formalisms of different kinds – VDM (based on the model-oriented approach), $Maude$ (based on the rewriting logic approach), and basic CCS^{dp} (based on the behavioural approach) – to produce representations of the case study, and to evaluate how well the different representations can be verified formally to determine whether or not they meet the termination requirement. We have deliberately chosen **not** to use workflow-specific formalisms for two reasons. First, because of our lack of fluency in them; and second, because we believe the models should be produced using general purpose formalisms.

IV. Conclusions

Our research is ongoing. Initial results demonstrate that in all three formalisms: functional interference between concurrently executing tasks is expressed in terms of interleaved actions, but not preemptive actions; unplanned dynamic reconfiguration can be represented; and the termination requirement can be verified. Basic CCS^{dp} expresses reconfiguration declaratively, and (therefore) very simply, which is useful for experimenting with different reconfiguration paths between configurations; but there is no support for how reconfiguration can be implemented. Verification is possible through model checking. In contrast, VDM expresses reconfiguration imperatively, and (therefore) in terms of mechanisms, which is useful for implementing reconfiguration; but complicates experimenting with different reconfiguration paths by including unnecessary detail. In contrast, $Maude$ expresses actions as rewrite rules, and reconfiguration as rewrite rules of rewrite rules, which is elegant; but it is not how a typical system designer thinks about reconfiguration.

Acknowledgements: The authors acknowledge the help given by numerous colleagues, in particular: Jeremy Bryans, John Fitzgerald, Regina Frei, Kohei Honda, Cliff Jones, Maciej Koutny, Richard Payne, Traian Florin Serbanuta, Giovanna Di Marzo Serugendo and Chris Woodford.

References

[1] F. Abouzaid, J. Mullins, M. Mazzara, and N. Dragoni. Modelling and analysis of dynamic reconfiguration in bp-calculus. In *Proceedings of the 6th KES International Conference on Agent and Multi-Agent Systems (KES-AMSTA 2012)*, volume 7327 of *Lecture Notes in Computer Science*, pages 534–543. Springer-Verlag, 2012.

[2] J. P. A. Almeida, M. Wegdam, M. van Sinderen, and L. Nieuwenhuis. Transparent dynamic reconfiguration for corba. In *Proceedings of the 3rd International Symposium on Distributed Objects and Applications*, pages 197–207, 2001.

[3] A. Bhattacharyya. *Formal Modelling and Analysis of Dynamic Reconfiguration of Dependable Systems*. PhD thesis, Newcastle University School of Computing Science, 2013.

[4] T. Bloom and M. Day. Reconfiguration and module replacement in argus: theory and practice. *Software Engineering Journal (Special Issue)*, 8(2):102–108, 1993.

[5] J. S. Bradbury, J. R. Cordy, J. Dingel, and M. Wermelinger. A survey of self-management in dynamic software architecture specifications. In *Proceedings of the 1st ACM SIGSOFT Workshop on Self-Managed Systems*, pages 28–33, 2004.

[6] C. Ellis, K. Keddara, and G. Rozenberg. Dynamic change within workflow systems. In *Proceedings of the Conference on Organizational Computing Systems*, pages 10–21. ACM, 1995.

[7] S. Fischmeister and K. Winkler. Non-blocking deterministic replacement of functionality, timing, and data-flow for hard real-time systems at runtime. In *Proceedings of the 17th Euromicro Conference on Real-Time Systems*, pages 106–114. IEEE Computer Society, 2005.

[8] G. J. Fohler. *Flexibility in Statically Scheduled Hard Real-Time Systems*. PhD thesis, Technische Universität Wien Institut für Technische Informatik, 1994.

[9] E. Fărcaş. *Scheduling Multi-Mode Real-Time Distributed Components*. PhD thesis, University of Salzburg Department of Computer Sciences, 2006.

[10] J. Kramer and J. Magee. The evolving philosophers problem: Dynamic change management. *IEEE Transactions on Software Engineering*, 16(11):1293–1306, 1990.

[11] M. Mazzara, F. Abouzaid, N. Dragoni, and A. Bhattacharyya. Toward design, modelling and analysis of dynamic workflow reconfiguration - a process algebra perspective. In *Proceedings of the 8th International Workshop on Web Services and Formal Method (WSFM 2011)*, volume 7176 of *Lecture Notes in Computer Science*, pages 64–78. Springer-Verlag, 2012.

[12] M. Mazzara and A. Bhattacharyya. On modelling and analysis of dynamic reconfiguration of dependable real-time systems. In *Proceedings of the 3rd International Conference on Dependability (DEPEND 2010)*, 2010.

[13] N. Medvidovic and R. N. Taylor. A classification and comparison framework for software architecture description languages. *IEEE Transactions on Software Engineering*, 26(1):70–93, 2000.

[14] J. Montgomery. A model for updating real-time applications. *Real-Time Systems*, 27(2):169–189, 2004.

[15] P. S. M. Pedro. *Schedulability of Mode Changes in Flexible Real-Time Distributed Systems*. PhD thesis, University of York Department of Computer Science, 1999.

[16] L. Sha, R. Rajkumar, J. Lehoczky, and K. Ramamritham. Mode change protocols for priority-driven preemptive scheduling. *The Journal of Real-Time Systems*, 1(3):243–264, 1989.

[17] D. B. Stewart, R. A. Volpe, and P. K. Khosla. Design of dynamically reconfigurable real-time software using port-based objects. *IEEE Transactions on Software Engineering*, 23(12):759–776, 1997.

[18] K. Tindell, A. Burns, and A. Wellings. Mode changes in priority preemptively scheduled systems. In *Proceedings of the IEEE Real-Time Systems Symposium*, pages 100–109, 1992.

[19] M. A. Wermelinger. *Specification of Software Architecture Reconfiguration*. PhD thesis, University of Lisbon Department of Informatics, 1999.